Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on Intel initiatives, technologies, platforms, and products.

Cover Story

Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

Featured Articles

Delivering in-depth reports on key platforms, products and technologies, our featured articles provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

Contact the Editor

To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

Archives

Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue.

Bookmarking

We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe* Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

Table of Contents

(Click on page number to jump to articles)

Note: Intel does not control the content on other company's Web sites or endorse other companies supplying products or services. Any links that take you off of Intel's Web site are provided for your convenience.

Cover Story

Extending and Expanding Moore's Law

Overview

Moore's Law—the prediction that the number of transistors on an integrated circuit (IC) will double every 18–24 months—has proven more accurate, lasted longer and produced more far-reaching changes than Gordon Moore could have anticipated when he made his prediction in a 1965 article in the 35th anniversary issue of *Electronics*. The law has since become virtually a blueprint for the electronics industry.

Now, rather than slowing down, the effect of Moore's Law is increasing. Intel expects to produce billion-transistor processors by the end of the decade, and the range of devices that can be manufactured in silicon is expanding. So, while processes described by Moore's Law have already transformed the world, the future impact is likely to be even more dramatic.

As silicon technology evolves, Moore's Law will catalyze the development of new application areas, bringing about the seamless integration of computing and communications and expanding the benefits of Moore's Law beyond today's digital realms.

Removing Roadblocks

Driving Moore's Law means reducing process geometries—shrinking the size of the devices populating the silicon. Intel researchers are aggressively identifying and eliminating any barriers that impede Intel's ability to extend Moore's Law. Intel breakthroughs in the past year alone have removed barriers to extending Moore's Law for at least another decade—and likely beyond.

Transistor size. Intel recently began production of a new 130-nanometer (nm) process that features 60 nm gate length transistors and six layers of copper interconnect. In June 2001, Intel announced it had developed transistors featuring structures that are just 20 nm in size. By the end of 2001, Intel demonstrated the world's smallest transistor, with a gate length of 15 nm. These incredibly small 15 nm transistors are approximately the size that will be needed in manufacturing in the year 2009**.

Heat and power dissipation. In Nov. 2001, Intel announced the development of a new type of transistor that helps create extremely power-efficient devices. Intel expects to incorporate elements of this new structure into its product line in the second half of this decade**.

Lithography. Extreme Ultraviolet (EUV) lithography is a technology breakthrough that allows the patterning of lines smaller than 50 nm. Intel leads a consortium of semiconductor companies, the EUV LLC (Limited Liability Corporation), that announced completion of the first full-scale prototype machine for making computer chips using EUV lithography. Intel anticipates building processors using EUV technology in the second half of the decade**.

Packaging. In October 2001, Intel announced an innovative packaging breakthrough, Bumpless Build-up Layer (BBUL), that grows the package around the silicon. This technique reduces the thickness of the package and enables the processor to run at a lower voltage. BBUL packaging technology is an example of an innovative way to embed powerful computers chips into very small spaces.

Wafer size. In conjunction with its new 0.13-micron (130 nm) technology, Intel has begun utilizing the latest 300-millimeter (12 inch) diameter silicon wafers. Producing chips using these 300-mm wafers will help cut costs by 30 percent and energy and water requirements by 40 percent.

New Applications for Moore's Law

The classic definition of Moore's Law, based on the doubling of transistor count, no longer fully reflects today's silicon opportunities. Moore's Law is expanding to encompass the emergence of exciting new technologies that take silicon beyond transistors.

The expanded Moore's Law means that silicon-based capability will double every 24 months by increasing the *count, complexity, and convergence* of components and technologies integrated on a chip. This combination of three vectors—device count, complexity, and convergence—provides a richer set of resources for increasing capability and functionality, and for improving the flexibility with which it can be applied.

Device count on a chip remains fundamental to Moore's Law. The more devices on a chip, the more computation, memory, and functional integration can be accomplished.
Complexity provides a second vector of increasing capability and flexibility. Increasing the complexity of transistors and other devices leverages the doubling cycle for device count. Examples of increased complexity include constructing transistors with more than one gate or vertically stacking multiple layers of transistors.
Convergence brings together multiple functions and heterogeneous technologies onto a single chip. For example, Microelectromechanical Systems (MEMS) make it possible to integrate mechanical and electrical components onto other components, thereby expanding their functionality.

More Speed, More Capabilities

The added dimensions of complexity and convergence offer increasing opportunities to make the benefits of Moore's Law available to a broader range of applications and products. The goal in each case is to not simply improve performance, but to deliver new functionality. By applying the power of Moore's Law to new classes of functionality, Intel is bringing about a new computing and communications geography. This expansion of Moore's Law makes these new applications more affordable and widespread, and opens the door to new areas of innovation. Here are three examples:

Wireless, self-organizing sensor networks combine silicon advances driven by Moore's Law with communications network research to enable thousands of small, embedded sensing devices to wirelessly connect. These ad hoc wireless sensor networks lend themselves to a wide range of novel applications, and can enhance productivity, safety, awareness, and efficiency.
Silicon radios will help extend expensive wireless chipsets into ubiquitous (and virtually free) silicon radios-on-a-chip. Eventually, Moore's Law will make it possible to have multiple silicon radios, with intelligent roaming, that take up little space on a chip, but add enormous capability and flexibility. This will make many critical activities function better, from running a supply chain to finding a lost child.
Silicon photonics brings optical networking technologies into silicon, extending the cost, size, and performance advantages of Moore's Law to a new and important arena. Intel recently demonstrated a prototype of a silicon-based tunable filter that puts the functionality of a $10,000 optical filter onto a $1 chip, advancing toward a goal of pennies apiece. This approach promises to make optical networking practical and affordable, not simply for the long-haul markets and wide area networks, but also for smaller networks and ultimately for chip-to-chip interconnects.

More Info

Read the full document on Expanding Moore's Law from which this article is based.
For more information on expanding Moore's Law and related technologies, please visit the following sites:
Intel Labs
Intel Research

Departments

Special Section: Expanding Moore's Law

Ad Hoc Sensor Networks: Proactive Computing Is Coming

Kevin C. Kahn
Intel Fellow and Co-Director of Communication and Interconnect Technology
Intel Labs
Intel Corporation

David E. Culler
Director
Intel Research Laboratory, Berkeley, California
Intel Corporation

Overview

Today, hundreds of millions of computing desktops and laptops are connected via the Internet. Tomorrow, hundreds of billions of embedded chips and sensing devices will be integrated into everything from key chains to baby cribs—and they'll all have the ability to compute, sense, and communicate.

At Intel's research facilities and elsewhere, the silicon advances driven by Moore's Law (simply stated, semiconductor capacity doubles roughly every 18 months) are being combined with communications network research to enable thousands of small, embedded sensing devices to wirelessly connect and use virtually no power. The result is a new computing paradigm enabled by these *ad hoc wireless sensor networks.*

The Age of Proactive Computing

Ad hoc wireless sensor networks represent a core set of functionality needed to move into this next computing era when computers will be directly connected to the physical world and begin anticipating what you might want to do next—sometimes even taking action on your behalf.

This new paradigm not only enhances productivity, but also can enhance safety, awareness, and efficiency. Imagine a crib that listens to a baby's breathing, a swimming pool that warns you if something falls in, a bracelet that informs you of an elderly parent's current health, or smoke detectors that help guide emergency personnel—all connected and communicating critical data where it needs to go.

In the everyday world, ad hoc sensor networking will help enable electronic devices to recognize each other. Motors will tell you when they need maintenance and your laptop will know when it's been dropped and needs a repair.

The potential uses extend far beyond the home. Researchers are exploring the use of tiny devices that can be dropped by airplanes over a large forest to find a lost child, or monitor the temperature and movement of a major forest fire. Crop conditions will be monitored down to individual plants. Sensor networks may also enable intelligent sensing of environmental conditions in homes and offices. Climate control could be more efficiently targeted to specific rooms, potentially reducing overall power consumption—a plus for the environment.

The role of humans will change in this new paradigm. Computers become more proactive, anticipating human needs and meeting them. The human will be at the pinnacle and in control, instead of in the middle, shuttling the information between the real world and the computer.

Technology Challenges

There are challenges to making ad hoc sensor networks a practical reality. These sensing devices, called motes, are actually wireless, embedded platforms that combine sensing, communication and computation. The computers must not only be able to sense the world around them, but also to act upon it. And they need to be networked.

The proliferation of silicon-based sensing devices could create a hundredfold increase beyond the growth of the Internet. There are a number of networking areas being researched today, including the creation of large, self-organized networks, making the networks delay-tolerant for applications such as deep space or rapidly moving sensors, and abstracting the myriad of small-networked devices to make them accessible to computing applications.

Moore's Law and its continued silicon technology improvements are helping make sensor devices smaller and less expensive and reduce their power requirements. However, many core challenges facing these new networks are "systems" challenges. For example, these new types of networks need to be easy to use and able to function without demanding human attention.

Intel Research Focus

The Intel Berkeley Research laboratory, located near University of California, Berkeley, is leading research efforts, working with the academic community and the industry to address many of these systems-level challenges. The Lab's research focuses on the architectural elements around these tiny, low-power wireless embedded platforms. This research also encompasses:

The need for a flexible, open operating system.
The networking technology challenges for large self-organized networks of sensor devices.
Higher-level services needed to make it easy to develop robust applications for ad hoc sensor networks.

The Lab's researchers have already created a prototype research platform that's enabling researchers to explore ways to use an ad hoc sensor net. The current platform measures 1-inch square, but in the future you can expect to see sensors only a cubic millimeter in size (see Figure 1), and eventually as small as a speck of dust.

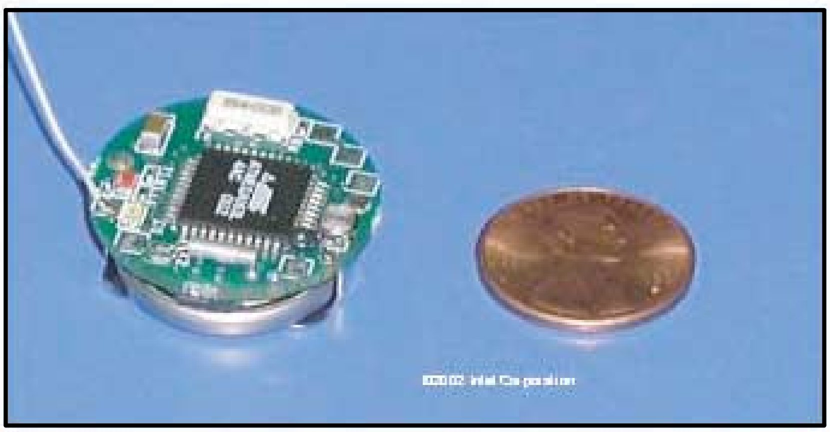

Figure 1. Advances in silicon technology will help drive sensor devices down to a cubic millimeter and smaller. Pictured above is a "dot-mote" currently in use and almost the size of a penny.

The ability to easily and cheaply deploy large, self-organized networks of sensor devices is only useful if these devices can sense useful data. Intel is also working on ways to connect devices to the physical world, which means developing a whole new suite of sensors and actuators: sensors that can sense what's going on around them, actuators that allow the devices and computers to change the world around them, and in a size that will fit on these minute device platforms. Intel's Microelectromechanical Systems (MEMS) research is just one example of the new kinds of sensors being developed, and Intel is applying its precision silicon capability to drive this effort.

Intel researchers are also working on precision biology—biochips. In addition to sensing dry, solid-state phenomena, our researchers are exploring the ability to sense wet substances, including biological materials and organic chemistry. This capability opens up every aspect of wetware, from health and pharmaceuticals to chemicals and refineries, and creates another new frontier for computing applications.

Summary

The new computing paradigm enabled by ad hoc wireless sensor networks will be key in making computing more proactive. Silicon-based sensors and ad hoc sensor networks represent exciting new technologies with broad societal impacts and a wide range of new commercial opportunities, and Intel is helping break down barriers that prevent these technologies from becoming a reality. Intel-led research is addressing many of the challenges of ad hoc sensor networks, and advances in silicon technology, in line with Moore's Law, are reducing the size, cost, and power of sensor devices.

More Info

Read the full white paper on Ad Hoc Sensor Networks from which this article is based. For more information on this and related technologies, please visit the following sites:
Intel Labs
Intel Research
Intel Research laboratory at Berkeley
Papers by David Culler

Author Bios

David E. Culler, professor in computer science at UC Berkeley, has taken a leave of absence to work for Intel as the university director of the Intel Research lab at Berkeley. His research is in areas of parallel architectures, programming languages and operating systems. Dr. Culler received his B.A. in mathematics from UC Berkeley, followed by an M.S. in computer science, and his Ph.D. in computer science from Massachusetts Institute of Technology.

Kevin C. Kahn is an Intel Fellow in the Corporate Technology Group and co-director of Communications and Interconnect Technology at Intel Labs. His research focuses on developing physical layer technologies for communications systems and system interconnects using RF, copper, and optical media. Dr. Kahn received his Ph.D. and M.S. in computer science from Purdue, and his B.S. in mathematics from Manhattan College.

MEMS: Expanding Functionality with Miniature Silicon Devices

Valluri R. Rao
Intel Fellow and Director of Analytical and Microsystems Technologies
Technology and Manufacturing Group
Intel Corporation

Overview

Microelectromechanical systems (MEMS) are mechanical structures fabricated on silicon, using semiconductor-compatible manufacturing techniques. MEMS are making it possible to integrate mechanical—as well as electrical—components onto chips.

By combining MEMS with the transistor scaling that Moore's Law provides, entirely new Microsystems can be created. These make it possible to squeeze more functionality into a smaller footprint. In essence, we're starting to expand the envelope of what is considered to be the processor to include all the pieces in the chain going from the processor to the outside analog world.

Expanded Functionality

MEMS are microchips whose functionality includes mechanical elements that can move. This motion is driven by voltages and currents, which can be brought into the silicon chip. As part of a closed-loop system comprising MEMS sensors and actuators and digital and analog processing chips, the MEMS devices can sense environmental changes and actuate based on these changes with the closed-loop control provided by the electronics. This creates new ways for computers to interact with the physical world.

As the MEMS Clearinghouse puts it, if microprocessors are the brains of the computer, MEMS adds eyes and hands, enabling the creation of microsystems that can sense and act on the environment. Potential uses are as wide-ranging as inertial sensors, microscopic positioning stages, micromirrors, microfluidic devices to cool high-power chips, biochips that detect hazardous chemicals, ultra-high-density storage devices, and many more.

Technology Challenges

MEMS technology enables a new generation of microminiature mechanical devices to be fabricated using process technologies similar to those used in the semiconductor industry. These mechanical structures may form systems that have complex functionalities, such as movable mirrors on a chip for projection displays and optical switching, or gyroscopes that monitor the motion of an automobile. Examples of a MEMS micromirror and a MEMS inertial device (gyroscope) are illustrated in Figures 1 and 2.

The term MEMS has traditionally been used in the United States to embrace everything related to sensors and actuators built with silicon-like manufacturing technology, as well as the end systems that incorporate sensors, processors, and actuators.

In Europe, MEMS are known as *microsystems* and in Japan as *micromachines*. For the purpose of this overview, we will use MEMS and microsystems interchangeably, although there are good reasons to classify MEMS purely as a manufacturing process (just like CMOS) used for fabricating the devices, and microsystems to describe the end systems that perform the task of sensing, computation, and actuation.

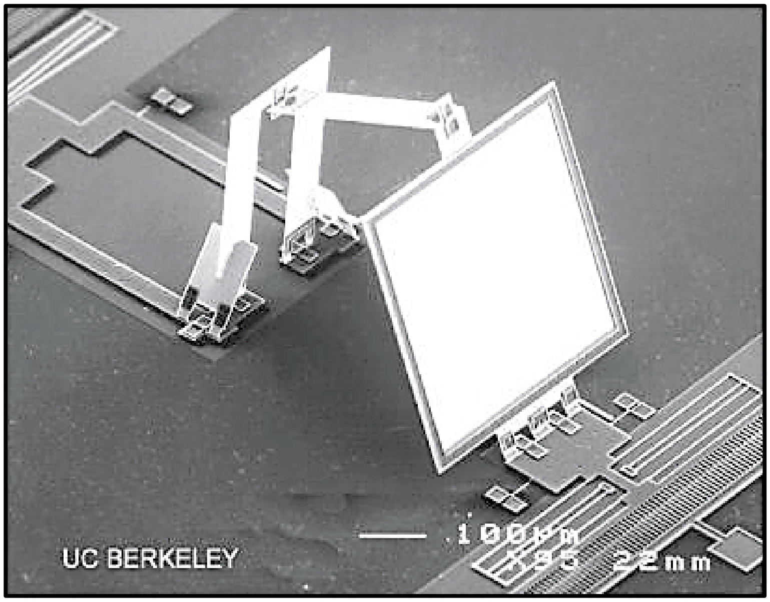

Figure 1. Micro Mirror
Courtesy: Berkeley Sensors and Actuators (BSAC)

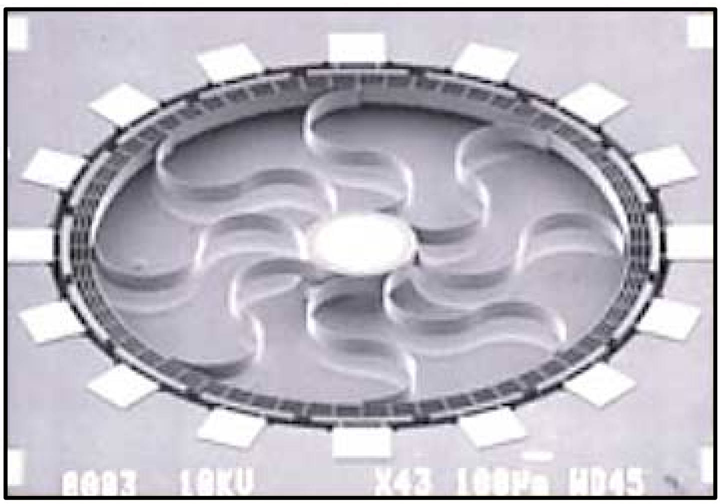

Figure 2. Vibratory Gyroscope
Courtesy: University of Michigan

Smaller Equals Faster

Intel's recent demonstration of a "TeraHertz" transistor capable of switching *one trillion times each second* provides the latest glimpse of where Moore's Law is headed. This "TeraHertz" transistor is expected to branch into mainstream electronic components that can run at 20 GHz and will contain over a billion transistors—within six years**.

The industry is advancing to decrease the element feature size of these components to 0.1 micron in the near future… then to 0.07 micron, 0.05 micron, and finally to 0.03 micron in the 2010** time frame. Intel has already announced plans for a 0.02-micron gate-length transistor to be incorporated in 0.05-micron processors and other components by 2007**.

At these size scales, the mechanical components fabricated on chips become so small that they can be made to move very quickly—*to vibrate*—in the multi-GHz range. This is fast enough to allow MEMS devices to serve as mechanical

radio components. The industry is on track to reap extra benefit from Moore's Law by applying the miniaturization trend to the mechanical as well as the electrical world, creating benefits far beyond the traditional metric of transistor count.

Intel Research Focus

Intel is pursuing MEMS research and development both internally and with academic institutions. Intel participates in a number of MEMS consortia, including ones at the University of California, Berkeley, and the University of Michigan. Intel is also making strategic investments in companies that are key innovators in MEMS technology. Intel is a founding member of the MEMS Industry Group, a trade association for the U.S. MEMS industry.

MEMS are leading candidates to address the demands of the wireless telecommunications ("RF") space. MEMS technology enables flexible wireless building blocks that can bring higher performance, smaller footprints, and tighter integration to RF-passive components used in wireless systems.

In addition to the passive electronics, MEMS technology has been applied to fabricate very-low-power bi-stable color displays, very sensitive directional microphones, and smart antennae. This blending of electronics with ultra-small mechanical moving parts creates a Moore's Law path that improves integration and reduces cost for the entire RF system.

You can find out more about the research efforts Intel is conducting on MEMS technology by linking to the white paper listed in the More Info section of this article.

Summary

The MEMS industry is on track to make the miniaturization benefits predicted by Moore's Law apply to the mechanical—as well as the electrical—world, creating advantages far beyond the traditional metric of transistor count. Because MEMS technology is based on many of the same processing techniques as conventional silicon, it is expected to explode into a very broad and diverse set of new applications—many of which have yet to emerge.

Intel, through its own research and engineering efforts, its support for university research, and its investment in innovative companies, is applying its leadership position in the silicon industry to the development of MEMS technologies.

More Info

Read the full white paper on Microelectromechanical Systems (MEMS) from which this article is based. For more information on MEMS and related technologies, please visit the following sites:
Intel Labs
Intel Research on MEMS
MEMS Clearinghouse (Information and Discussion)
MEMS Industry Group

Author Bio

Valluri R. Rao is an Intel Fellow in the Technology and Manufacturing Group at Intel. Dr. Rao is responsible for directing Intel's MEMS research and development activities. He previously guided the development of advanced analytical tools and methods for microprocessor performance characterization, silicon debug, and yield enhancement. Rao received B.A., M.A., and Ph.D. degrees in electrical engineering from Jesus College, Cambridge University, where he also was a post-doctoral research fellow until joining Intel.

Silicon Photonics: New Opportunities for Silicon

Shekhar Y. Borkar
Intel Fellow and Director of Circuit Research
Intel Labs
Intel Corporation

Mario Paniccia
Principal Engineer and Director of Photonics Research and Development
Intel Labs
Intel Corporation

Overview

Silicon photonics aims to determine how to use silicon and standard silicon processing techniques to build optical devices. The concept is based on developing optical building blocks that give active functionality, rather than simple, passive optical wave-guiding. In the future, these tiny silicon building blocks can be selectively placed into optical modules, reducing cost and size.

The core of Intel's silicon photonics research is based on a novel approach that allows dynamic electronic control of optical signals with no moving parts. Some of the devices that could result from this research include optical filters, fast switches (~10 ns), and very fast optical modulators (>GHz). To date this effort has produced a number of functional optical devices—all in silicon.

Technology Challenges

Today, both home and business computers are limited less by processor performance than by the rate at which data can be transmitted between the processor and the outside world. Major corporations, financial institutions, and virtually all businesses and consumers demand instant and reliable transmissions, whether their data is traveling across the street or around the world. Corporate LAN and Internet traffic already exceeds telephony traffic, and Internet traffic has been doubling every year.

Copper-based networking can no longer keep up, so the telecommunications industry has turned to fiber optics to fulfill the growth demands of Internet traffic. But optical networks are arcane and expensive. Because of their high costs, their use has been limited primarily to long-haul and backbone networks.

Dense Wavelength Division Multiplexing (DWDM) is a technology that boosts capacity on fiber-optic cables by sending multiple wavelengths of light on a single fiber. Since the first DWDM systems were deployed in the mid-1990s, the number of channels supported on a fiber has increased from less than 16 to more than 40. Multiple channels are now transmitted over longer distances, operating at 10 gigabits per second (Gbps).

The evolution toward faster data rates will drive the fiber-optic industry to move next to 40 Gbps, and to even higher data rates in the future. With the combination of higher data rates and DWDM capabilities, telecommunication companies will be able to transmit a trillion bits of data per second on a single fiber — a rate that would exceed the total traffic on the entire Internet today.

Intel Research Focus

Intel expects to allocate over $4 billion dollars in research and development in 2002. Intel Labs is focusing resources in four key areas: silicon and manufacturing, compelling computing architecture, communications, and new Internet capabilities. Photonics will play a role in all of these areas.

Intel has many active internal research efforts focused on photonics. The Intel Labs research community is continuously creating new materials, new structures, and new architectures to enable the evolution of photonics from discrete devices to integrated photonic systems. These efforts range across a wide spectrum of communication areas, progressing from 10-Gb transmitters and modules, to high speed ICs, to more longer range research activities looking at chip-to-chip interconnects. Intel also teams with academic researchers and supports innovative university research projects.

Intel Labs is conducting research in silicon photonics, developing ways to use silicon and standard silicon processing

techniques to build optical devices. Some of the devices that could result from this research include optical filters, fast switches (~10 ns), and very fast optical modulators (>GHz). Silicon photonics process development activities are being conducted in Intel's facility in Israel, Fab 8 (see Figure 1), where our MEMS (Microelectromechanical Systems) activities are also being developed.

One silicon photonics device that Intel is working on is a tunable optical filter that we showcased at the Intel Developer Forum Conference Spring 2002. The filter is just a few microns wide by a couple of millimeters in length and can filter wavelengths in the DWDM spectrum.

The small form factor outlines the potential value of developing these photonic devices in silicon. However, there is a fundamental limit that stems from the inherent nature of photons, in shrinking these devices to dimensions smaller than a few microns. Producing optical devices in silicon affords the opportunity to integrate various functions together in a much smaller form factor than exists today.

In addition, the industry can use silicon combined with standard IC manufacturing and assembly technologies to produce novel, low-cost packaging and assembly technologies. This is the real opportunity and value of silicon photonics, irrespective of specific device or applications.

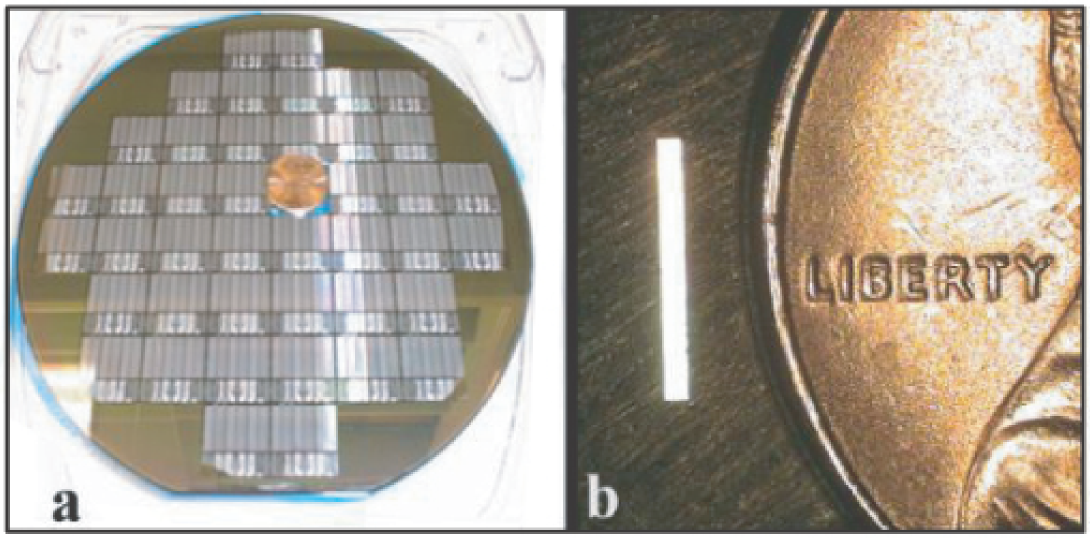

Figure 1. (a) Silicon photonics wafer processed at Fab 8 in Israel.
(b) Three optical filters fit into this tiny piece of silicon.

Summary

The Internet explosion has spurred the growth of the photonics industry. The photonics technology revolution promises high-speed, high-capacity fiber-optic Internet communications for use in next-generation applications in education, medicine, entertainment, and commerce.

Intel is actively developing solutions that enable high-volume manufacturing, integration and more cost-effective packaging capabilities. The goal of bringing optical networking down to the curb and ultimately to the PC can only be met if high-performance, low-cost, low-power optical components are available. Intel Labs will continue to research approaches to span smaller distances with photonic interconnects between components, and ultimately between circuits.

In some ways, the photonics industry is where the semiconductor industry was in 1970. However, 30 years ago, we didn't have the Internet or the high-volume fabrication plants that we have today. We also did not have the expertise to increase microprocessor capacity, garnishing another 25 MHz that we now measure in months instead of years. With photonics, circuit densities can continue to double every 18 months for decades upon decades to come. That is expanding Moore's Law.

More Info

Read the full white paper on silicon photonics from which this article is based. For more information on photonics and related technologies, please visit the following sites:
Intel Labs
Intel Research

Author Bios

Shekhar Y. Borkar is an Intel Fellow in the Enterprise Platforms Group, and director of Circuit Research at Intel Labs. Borkar is responsible for directing research in low-power circuits and high-speed signaling for Intel's future microprocessors. He received an M.S. in electrical engineering from the University of Notre Dame, and M.S. and B.S. degrees in physics from the University of Bombay.

Mario Paniccia is principal engineer and director of Photonics Research and Development for Intel. Dr. Paniccia has been involved in many areas of optical research throughout his career, including developing optical testing technology, and working on optical interconnects, optical clocking, and optical communications. He leads a research team focused on developing silicon-based photonic devices using standard silicon processing techniques. He received his B.S. in physics from State University of New York at Binghamton, and his Ph.D. in physics from Purdue University. Dr. Paniccia has published numerous papers, been awarded 15 patents, and has 25 additional patents pending.

CMOS Radio: Ubiquitous, Silicon-Based Wireless Connectivity

Steve Pawlowski
Intel Fellow and Co-Director of Communication and Interconnect Technology
Intel Labs
Intel Corporation

Overview

Before the advent of digital processing, radios were designed entirely of analog circuitry. As advances in the cost and scale of CMOS technology provided digital processing power at a reasonable price, digital signal processing (DSP) began to play a major role in overall communication system designs. Ever-improving DSP techniques have enabled improvements in communications, as reflected by Moore's Law.

Today, we are experiencing the power of DSP techniques through many wireless radio frequency (RF) communication applications. Wireless wide area networks (WWANs, or cell phones), Wireless Local Area Networks (WLANs), and Wireless Personal Area Networks (WPANs) all employ sophisticated communication techniques. Some of these techniques include complex modulation schemes, powerful new error-correcting codes, decoding algorithms to combat the effects of channel fading, and so on.

All these techniques are being enabled, cost-effectively, by the increasing capabilities of digital processing, as predicted by Moore's Law. At the same time, CMOS technology has enabled digital devices to be produced in high volumes, again in a cost-effective manner, enabling larger markets.

New CMOS Capabilities

Until recently, high-frequency wireless communications applications have used technology processes such as Gallium-Arsenide (GaAs) to obtain the performance needed from the RF Analog Front End (AFE) circuits. Although these processes provide the functional performance required by radios today, they do not support the cost/scalability business model of standard CMOS that is reflected by Moore's Law.

The higher switching speeds that result from the smaller geometries being developed in CMOS are enabling the design of analog circuits at very high frequencies, with very good gain and linearity. It is this new frontier where Intel Labs is focusing research on the design of analog RF circuitry, which utilizes the same CMOS process technology and device set that Intel uses to manufacture its microprocessors and chipsets.

This new capability will allow analog circuit designs to scale with the digital capabilities predicted by Moore's Law. Analog solutions implemented in CMOS will achieve high performance, functionality, and bandwidth while maintaining low cost, small size, high quality, and robust architecture across the wireless market (see Figure 1).

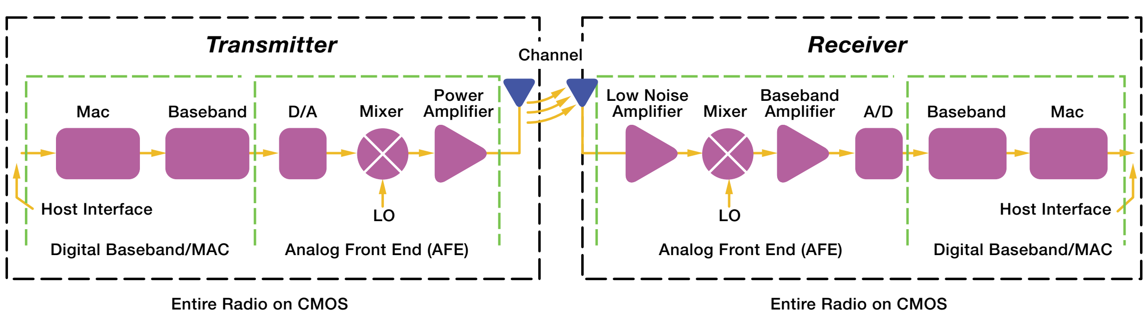

Figure 1. Essential Components of a Wireless Device

Intel Research Focus

Intel Labs research is addressing the challenges involved in utilizing a process technology that is optimized for digital switching transistors, at ever decreasing supply voltages, and synthesizing analog RF circuits. This is a leading area of industry and academic research, and these developments will pave the way for higher degrees of integration for radio chipsets with the attendant cost/scalability benefits.

To aid in the development and rapid deployment of new radio designs, Intel Labs is defining a reprogrammable, reconfigurable, digital communications platform. This development platform will use high-powered processors, DSP processing techniques, and high-speed digital logic to support high data rates (in excess of 100 Mbps) and will be reconfigurable to execute multiple wireless protocols.

For example, it will be scalable to enable 802.11a operation or an experimental Ultra Wide Band (UWB) prototype application operating at speeds greater than 100 Mbps. UWB is an ingredient technology for high-bandwidth wireless communications between devices at very low ranges. Intel Labs is investigating this technology to understand its behavior and potential applications.

In support of the reconfigurable digital communications platform, Intel Labs is researching a reconfigurable micro-coded accelerator to be used in developing new signal processing algorithms and techniques. The reconfigurable accelerator will use precisely defined computational elements, connected in a proposed hybrid-MESH network topology that can be programmed to execute different communications solutions at an energy-per-computation cost less than that of current DSP solutions.

This type of architecture will enable the realization of systems supporting different protocols, including improved capacity, robustness, and flexibility, while requiring much less energy for computation.

The research firm Cahners* In-Stat Group predicts that by 2005 at least 50 percent of over 900 million cellular phones will be data-enabled through either WPAN or WWAN next-generation radios. In addition, existing WLAN technologies continue to emerge to support applications with both personal and handheld computers. Users are sending and receiving data content over all these wireless protocols.

Today, they must carry different hardware solutions for each wireless frequency or technology, and in many cases must reconfigure their machines to access data using the required protocol. As we move into the future, the entire radio has the potential of being integrated onto a single chip. Eventually, users will have a wireless subsystem that employs multiple AFEs at multiple frequencies supporting multiple protocols (WWAN, WLAN, and WPAN).

Summary

Intel is developing the "radio of tomorrow." The Intel vision will be achieved when all the RF functionality becomes part of an existing component. For notebooks, the "component" may be the existing chipset. For smaller computing devices it may even be the processor.

As Intel applies its silicon technologies across a new geography of computing and communications paradigms, a more integrated continuum of computing and communications services is made possible. At the end of the day, it all comes back to the enduring power of Moore's Law and Intel's desire to abide by the law.

More Info

Read the full white paper on CMOS Radio from which this article is based. For more information on this and related technologies, please visit the following sites:
Intel Labs
Intel Research

Author Bio

Stephen S. Pawlowski is an Intel Fellow in the Corporate Technology Group and co-director of Communications and Interconnect Technology at Intel Labs. Pawlowski leads the research and development of wireless (RF), copper, and optical interconnect technologies for Intel's platforms and systems, as well as the definition and implementation of systems technology and ingredient roadmaps across Intel's product groups. He graduated from the Oregon Institute of Technology with a joint degree in electrical engineering technology and computer systems engineering technology, and obtained a master's degree in computer science and engineering from the Oregon Graduate Institute.

Desktop

Cruising with PC 'Hot-Rods' at IDF

Gabe Achanzar
Tech Leadership Program Manager
Desktop Platforms Marketing Group
Intel Corporation

Overview

Expect the unexpected. Because, as Intel found out, that's what you get when you let some of the planet's most innovative hard-core PC enthusiasts get creative with the latest Intel® Pentium® 4 processor.

You get a PC that looks like it should be home to a school of guppies, but is definitely 'Not-a-Fishtank.' You also get a PC with a case that at first glance appears to be empty, another with an industrial case design that hearkens back to the era of the U.S.-Soviet space race, a retro-look, 1940's sci-fi PC clad in copper with a built-in cigarette lighter, and a see-through neon-trimmed "X"-shaped PC with electroluminescent highlights.

These innovative PC designs are examples of the new competitive sport of "case modding" promoted by the Cyberathlete Professional League* (CPL). Similar to hot-rodders who customize car bodies, PC "modders" toss out the rules to stretch the boundary between PC technology and art.

The winners of the CPL's 2001 Computer Case Contest held during the December Winter Championships were on display at the Intel Developer Forum (IDF) Conference Spring 2002 in San Francisco. All feature the Intel Pentium 4 processor, and as you are about to see, these are definitely not your ordinary PCs.

"Not-a-Fishtank PC"


Figure 1. "Not-a-Fishtank PC" Modder: Troy Ervin

Any resemblance between this PC and an actual fish tank is purely circumstantial (see Figure 1). Modder Troy Irvin, whose day job is computer engineering, used a cardboard spool to build the chassis, which conceals a stealth CD-ROM. The machine also includes a modified 400-watt power supply, a 300-degree transparent plexiglass viewing area, a see-through top, and a red LCD text display. A pair of neon sticks illuminate the PC's interior.

Known as "T-Bone" within the PC modding community, Troy is a serious gamer as well as a PC modder. He is also a vegetarian, skateboarder, and in-line skater, and he coaches Little League baseball and youth football.

"Invisible PC"

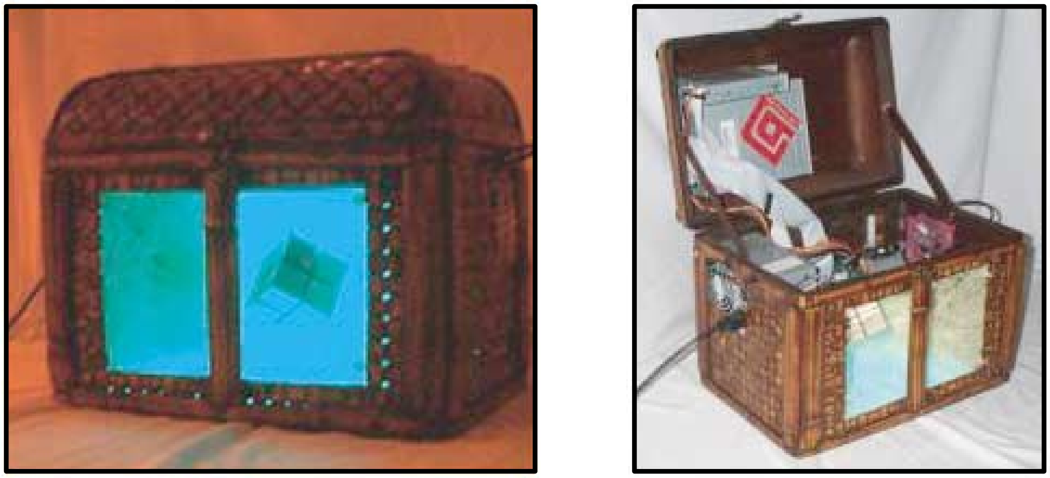

Figure 2. "Invisible PC" Modder: Dean Liou

At first glance this low-tech woven basket appears empty, but in reality the small-footprint case encloses a powerful Pentium 4 processor-based PC (see Figure 2). The optical illusion results from a mirror mounted at a 45-degree angle to the ‘floor’ of the case, and the effect is enhanced by the glow of electroluminescent lighting. The basket flips up to reveal the PC motherboard and stealth CD-ROM.

Dean “Envador” Liou is a prolific modder whose work is sought-after by collectors. He is known throughout the modding community for creatively merging everyday objects with PC technology with creations including the “barbecue PC,” the “toilet PC,” and the “PVC PC.” Dean is a network consultant who sees PC modding as a form of artistic expression.

"Supreme Soviet"

Figure 3. "Supreme Soviet" Modder: Anthony Stanley

The Supreme Soviet features the unmistakable symbols of the old Soviet Union (see Figure 3). The hammer and sickle

is backlit by red neon, and design elements on the case are accented in yellow and red electroluminescent trim. The PC features an LCD display that allows the user to input and view video, and a laser light show beams through the top cutouts of the case. Designer Anthony Stanley explains that the red star on the front of the case also functions as the power button, constructed from two plastic discs, a momentary switch, five screws, and some ballpoint pen springs. The disc allows light to shine through from the rear when the power is on, turning the star red. With the power off, the star blends in with the case metal.

Anthony "[H]amhouke" Stanley is the winner of the December 2001 CPL Case Modding Competition. Anthony is an Information Security Specialist, full-time college student, and adult leader in the Boy Scouts of America. His current projects include a commission from the *Official Xbox* Magazine* and the launch of a new modder's magazine, *InovaPC*.

"CopperClad Case"

Figure 4. "CopperClad Case" Modder: Barry Collins

Modder Barry Collins thought it would be a novel idea to create a PC computer case out of copper instead of the usual aluminum (see Figure 4). The resulting PC has a retro look reminiscent of the computers featured in classic science fiction dramas. The PC has neon accents and features the added touch of a working lighter. The front-panel LED satisfies the need of many tech enthusiasts to monitor internal case temperatures, but the innovations don't stop there. The case also comes with adjustable fans that provide the optimum cooling when required, and can be turned-down to quiet the unit.

Barry "Cold Dog" Collins is a research technician, an avid gamer, and a bona-fide technology enthusiast who attends computer and technology shows whenever he can find the time. He rounds out this techie side with his interest in Siberian Husky dogs, and travels throughout the south to American Kennel Club* dog shows.

"X-Gate"

Figure 5. "X-Gate" Modder: Mark Weitz

Designer Mark Weitz describes X-Gate as a "geezer's way" of showing his preferred gaming platform, featuring a Pentium 4 processor for exciting experiences matched by a cool design (see Figure 5). As Mark says, "The Pentium 4 processor gives me the processing power I need when editing film, and without a doubt it provides the extra speed I need when playing younger players." Highlights of the X-Gate include a water cooling system, internal electroluminescent highlights, and neon trimmings visible through a hand-cut fluorescent plexiglass panel. Mark says, "Modding to me is about art and function. I don't want to hide my box. I want people to see it."

Mark "geezer" Weitz works as a representative for a medical company, and his favorite pastimes include gaming, motorcycles, and digital photography. Everyone in his household is involved with computers. His wife Judy (Sparkle) runs Compucast.com, a New Orleans-based Internet consulting and Web site development firm. His oldest son, Matthew (Fractal), 22, is a computer science major at LSU and was the first to introduce geezer to the joys of modding. Younger son Michael (Macros), 17, is a nationally ranked gamer and substitutes for his dad whenever Mark finds himself getting "owned" by better players.

About the CPL

The Cyberathlete Professional League (CPL) is a leading professional computer gamer's league and conducts competitive multi-player tournaments worldwide. With members from over 50 countries, the CPL also holds computer case modding competitions, product exhibitions, and workshops for computer enthusiasts.

In the gaming community, processing power is a key ingredient, and so is reliability. These requirements help explain why the Intel Pentium 4 processor, Intel® D850GB motherboard, and Intel® Pro/100 S Desktop Adapter are recognized as "Official Computer Equipment" of the CPL.

The next C3 (Computer Case Contest) will be held during the CPL Summer Championship in Dallas, Texas, July 20–24, 2002.

Summary

To win one of the competitive case modding competitions sponsored by the Cyberathlete Professional League (CPL), your first step should be to throw out all the traditional rules for what a PC should look like.

Modders who took top honors in the CPL C3 2001 World Championships proved that a PC can resemble a fish tank, an industrial designed hold-over from the Soviet space program, a seemingly empty picnic basket, a retro-designed radio, or a stylized transparent letter "X" complete with neon lights.

What all of these designs have in common is a radical sense of design and innovative PC technology based on the power and reliability of the Intel Pentium 4 processor. In addition to providing “eye candy” for attendees at IDF Spring 2002, the winning designs provide a fresh look at how creating a PC can be fun for designers and viewers alike.

More Info

Visit the Cyberathlete Professional League site.
For info about the future of technology and its impact on the PC experience, visit the Intel Home Computing Web site.

Author Bio

Gabriel Achanzar is an admitted tech enthusiast who believes designing innovative PCs can also be a lot of fun. In his nearly two years at Intel, Gabe has put this philosophy into practice as manager of the Innovative PC Program and in his current role as manager of Intel’s Tech Leadership Program in the Desktop Platforms Marketing Group. With an M.B.A. in technology management, he has a patent pending in the area of ultrasonic transducers.

Initiatives & Technologies

Differentiating PCs in a 'Toaster World'

Robin Getz
Analog Devices

Overview

The days when PC OEMs (original equipment manufacturers) and system integrators were able to differentiate their product offerings based on memory, CPU speed, and hard drive size are gone. With the quality of today's Intel® reference designs, system designers are left with the option to add or subtract features like LAN and front-panel USB. Manufacturing and selling PCs is becoming more and more like manufacturing and selling toaster ovens in that product differentiation is being based on brand, service, design and color of the box, and so on.

This article will demonstrate how motherboard designers can add differentiating features to their designs, at little or no cost delta, allowing marketing and sales to differentiate on real-world functionality that consumers desire. The first feature that will be addressed is acoustic reduction through fan control.

New monitoring and control devices from Analog Devices and new functionality in the Intel® Pentium® 4 processor enable system OEMs to manufacture quieter PCs.

Thermal Management

The objective of thermal management is to ensure that the temperatures of all components in a system are maintained within functional limits. The functional temperature limit is the range within which the electrical circuits can be expected to meet their specified performance requirements. Operation outside the functional limit can degrade system performance, cause logic errors, or cause component and/or system damage. Temperatures exceeding the maximum operating limits may result in irreversible changes in the operating characteristics of the component.

However, it is important to understand that the thermal solution that provides the lowest temperature tends also to be the worst acoustically (fan speed or quantity), as well as the most costly (heat sink size). The optimal thermal solution will be just effective enough to ensure that the temperature limits of the system are met, while also minimizing the noise, all at the lowest possible cost.

The thermal monitor found in the latest Intel Pentium 4 processors allows system designers to design lower cost thermal solutions without compromising system integrity or reliability. By using a factory tuned, precision on-die thermal sensor and a fast acting thermal control circuit (TCC), the processor can keep its die temperature within factory specifications under nearly all conditions without the aid of any additional software or hardware. The thermal monitor thus allows the processor and system thermal solutions to be designed much closer to the power envelopes of real applications, instead of being designed to the much higher maximum processor power envelopes.

The thermal monitor controls the processor temperature by modulating (starting and stopping) the processor core clocks. The processor clocks are modulated when the TCC is activated. Once enabled, the TCC will activate only when the internal temperature is very near the temperature limits of the processor. When the temperature has returned to a non-critical level, TCC goes inactive and clock modulation ceases. This is a trip-point thermal sensor, which ensures that processor temperature in the hot spot does not exceed safe parameters. It is, in essence, a safety mechanism.

Processor performance will be decreased when the TCC is active; however, with a properly designed and characterized thermal solution, the TCC should never be activated, or will only be activated briefly during the most thermally intensive applications.

An external signal, PROCHOT# (processor hot) is asserted at any time that the temperature of the thermal monitor is above its internal limit. The temperature at which the thermal control circuit activates is not user configurable and is not software visible. Since each part is calibrated with respect to the maximum case temperature, if the maximum case temperature is never reached, the TCC will never become active.

In addition to the thermal monitor, all Intel® processors include an on-die thermal diode. An external thermal sensor such as the ADM1032, ADM1031 or ADM1027 produced by Analog Devices should be connected to the on-die processor thermal diode to monitor the die temperature of the processor. This thermal diode is separate from the thermal monitor's thermal sensor and cannot be used to predict the behavior of the thermal monitor. Because of the speed of temperature changes in "hot spots" and thermal gradients across the die, it is possible that the TCC may be active while the thermal diode is measuring a temperature within normal ranges.

While the heat generated on a specific part of the die is dissipated to the surrounding silicon as well as to the package, the inefficiency of heat transfer in silicon and between the die and the package results in temperature gradients across the surface of the die. Therefore, while one area of the die may have a temperature well below the design point, another area of the die may exceed the maximum temperature at which the design will function reliably. Figure 1 is an example of a simulated temperature plot of the Pentium 4 processor. The two arrows point to the center of the die, or the hottest thermal area, where the TCC is located, and the upper left where the thermal diode is located.

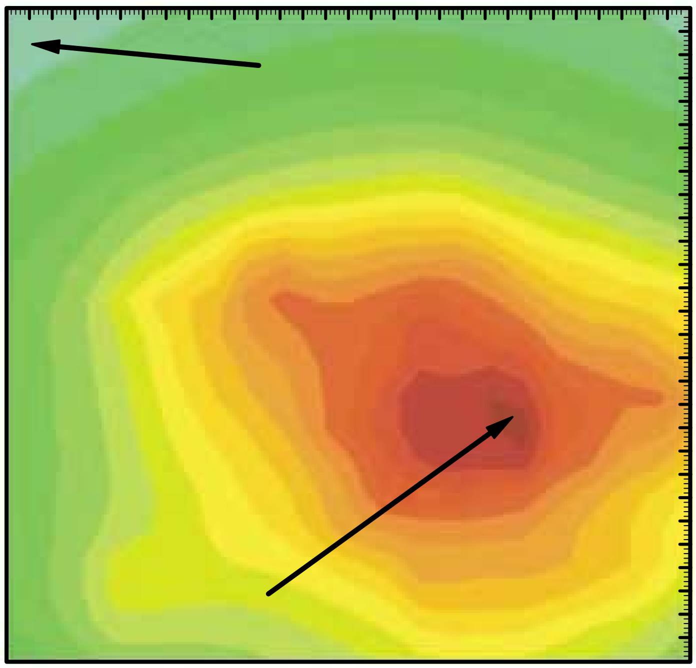

Figure 1. Simulated temperature plot of the Intel® Pentium® 4 Processor.

As a result of the cross-die temperature variations it is very important that the temperature detection mechanism (the integrated thermal sensor in the case of the Pentium 4 processor) be located at the hottest spot on the die. It is important to realize that because the TCC trip-point temperature is calibrated with respect to case temperature, when the TCC trips, the temperature of the thermal diode will vary with specific processor, heat sink solution, and thermal interface material. It is impossible to provide data with respect to the thermal diode, and correlate that to when or if the TCC is close to going off.

There are two main issues when dealing with processor thermals:
Ensuring that the thermals do not exceed functional limits. This is a function of the TCC, which is located in the hot spot, and can respond fast enough to handle the possible thermal transients.
Tracking overall temperature. This is the function of the thermal diode and the external thermal diode monitors. This monitoring function is away from the hot spots and is less likely to see high-speed thermal transients that the TCC protects the processor from. The thermal diodes can be used for automatic fan speed control to keep the system as quiet as possible, and to minimize the times that TCC events occur.

Importance of Noise

The problem with noise is not only that it is unwanted, but also that it negatively affects human health and well being. It has long been understood that a quiet work environment is a more productive one. In the not-so-recent past, devices that created moderate levels of noise (air conditioning, PCs, copy machines, and so forth) were thought to be acceptable. However, according to a study published in the *Journal of Applied Psychology* (Vol. 85, No. 5, pp. 779-783, Oct. 2000) ***Stress and Open Office Noise***, this is no longer true. These findings suggest that even moderately noisy open offices may contribute significantly to health problems such as heart disease (due to elevated levels of epinephrine, a stress hormone) and musculoskeletal problems.

Sound levels also impact our children's ability to learn. In a study by Gary Evans and Lorraine Maxwell at Cornell University (1997), it was found that children whose schools were affected by aircraft noise did not learn to read as well as those who were in quiet schools. The researchers compared children in a noisy school with similar students in a quiet school and found that children in the noisy school had difficulty acquiring speech recognition skills, impacting on the ability to learn to read.

To ensure that noise levels in classrooms are at acceptable levels, in 1999 the Access Board, which develops accessibility standards under the Americans with Disabilities Act (ADA), voted to collaborate with an existing Acoustical Society of America (ASA)/ANSI Working Group on Classroom Acoustics to develop recommendations for classroom acoustics. The group completed a final draft, which included a 35 dBA background noise specification, in January 2001, and submitted it for review and ratification to the ASA/ANSI Committee on Noise.

It is not just in the workplace and schools where end users are becoming more concerned with noise. The instruments of noise: television sets, computers, stereo systems, vacuum cleaners, and some toys, are plentiful in homes today. Limited studies have been conducted with younger children. One study in a residential setting found that 12-month-old infants in noisy homes exhibited less mastery-oriented play behavior with their toys than their counterparts in quieter homes.

Research findings suggest that exposure to uncontrollable noise may make children more vulnerable to learned helplessness. Learned helplessness means that individuals learn that the outcomes of their behavior are independent of their actions.

Most of this research has been with school-aged children, including kindergartners. One study found that children attending a noisy school were less likely to solve a challenging puzzle and to persist at it as well. Another study found that children exposed to noise were more likely to abdicate their choice for a reward to their teachers. The children decided to let the adult pick a prize for them rather than exercise their option to do so. Teachers in noisy schools also report greater difficulty in motivating children to do their schoolwork.

As more adults and children spend more and more time at personal computers, the ergonomic and environmental demands on these machines increases. Today there is an increasing awareness, both among PC users and manufacturers, that a low level of PC acoustic emissions must be regarded as an important factor in comfortable work and home environments. End users are unwilling to accept yesterday's solutions of over-design—putting more and more fans on a desktop computer, making it sound like a small airplane.

Fan Control

The real test is to implement automatic fan control mode and see if it does, in fact, affect processor thermals. In a simulated test, the combined effects of qJS (thermal resistance between sink and ambient) seemed to indicate that if the fans were not spinning quite as fast, and if the heat sink temperature was slightly higher at the beginning of the thermal event, it would not make a difference in the overall performance of the system.

This could be achieved when the heatsink fan speed is controlled via the on-die processor thermal diode. As processor temperature increases, fan speed will increase; as processor temperature decreases, fan speed will decrease. In the test, as fan speed decreased during the idle sections before and after the thermal event, the processor and heat sink were running marginally hotter than when the fan was running full on. However, even in the worst case, the processor temperature did not exceed normal operating parameters.

When using the AFC (automatic frequency control) methodology, during the power on self-test of the platform, the BIOS programs the temperature at which the fan turned on, the minimum fan duty cycle, the temperature range, and any temperature hysteresis, if desired. From a cool state, as temperature increased and approached T_{MIN}, the fan stayed off. Once the temperature reached T_{MIN}, the fan turned on to the minimum duty cycle.

As the temperature continued to increase, the fan speed also increased until it reached 100 percent at $T_{MIN} + T_{RANGE}$. As temperature decreased, the fan speed decreased until the temperature reached T_{MIN}. To keep from continually cycling on and off, the fan continued to run at the minimum speed until it reached the hysteresis point, where it turned off and waited until the temperature rose again. For design purposes, if the system designer does not want the fan to turn off at the hysteresis point, it is possible for the fan to run at the minimum operating point continually.

Since automatic fan control (AFC) does not adversely affect the thermal performance of the system, and it has been determined that processor temperature is the best input for fan control, then fan control based on processor temperature for acoustic betterment of the system can be explored.

Since each Pentium 4 processor is factory calibrated and individual systems will have different heat sinks, fans, chassis air flows, and amounts of heat generated, it is impossible to indicate the most appropriate T_{MIN} and T_{RANGE} for every system. Even for systems that have the same chassis and same internal components, the variability of processor, assembly, altitude, and other parameters make pre-defining a T_{MIN} very difficult. If the T_{MIN} is too low, excess acoustics are generated. If the T_{MIN} is too high, TCC clock modulation will occur, and performance will be impacted.

By taking advantage of the self-calibration features of the ADM1027, the system can tune its thermal response of that specific system and provide the quietest system possible. The addition of an operating point register allows the T_{MIN} to self-adjust. If the actual temperature exceeds the operating point, then T_{MIN} is reduced. As T_{MIN} decreases, the slope of the control line is constant, which increases the speed of the fan at the same temperature.

As shown in Figure 2, the ADM1027 has self-adjusted T_{MIN} to the new value t_{MIN} in order to keep the actual temperature below the operating point set at T_{OPER}.

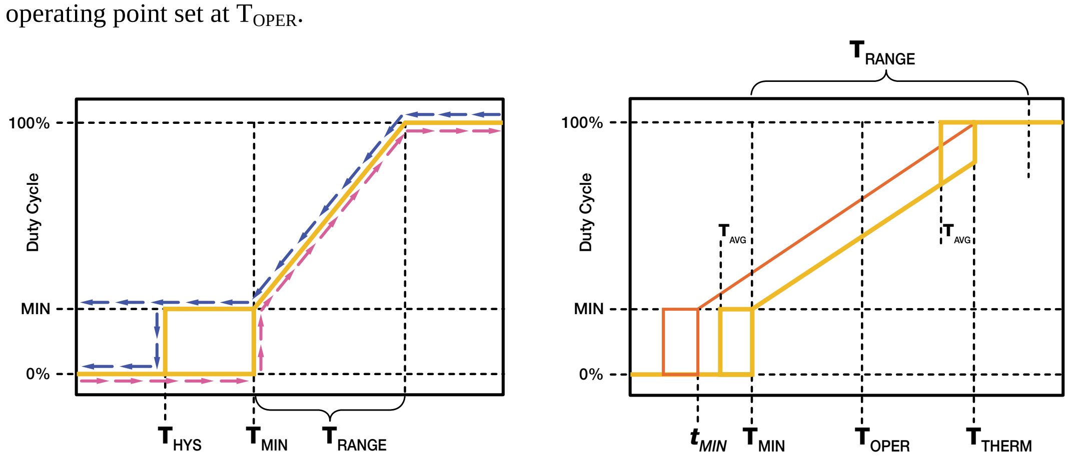

Figure 2. ADM1027 self adjusts T_{min} to the new value t_{min} to keep the actual temperature below the operating point set at T_{oper}

By using the ADM1027, which monitors the Pentium 4 processor's PROCHOT# pin, the ADM1027 can self-calibrate to just below the temperature at which the TCC becomes active. In this way, the temperature of the platforms will be the hottest possible, without running into thermal-related performance issues. As was stated previously, the hotter the system, the quieter the system.

The self-calibration functionality in the ADM1027 allows the system to become independent of fans variation, ambient air temperature, chassis design processor heat sink, thermal interface material (material between the processor and heat sink), altitude, networking cards, CD-ROMs, DVDs, hard drives, memory, chipsets, and other PCI cards and computer placement; all of which are creating or moving unknown amounts of heat in the system.

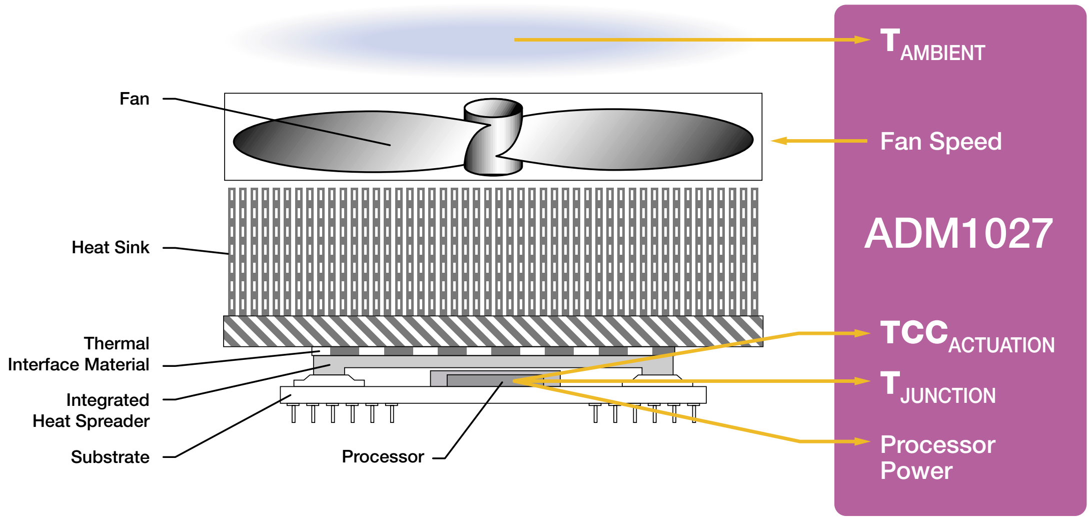

Figure 3. ADM1027 *dB*COOL* controller.

It is only with this type of self-calibrating functionality that the PC manufacturer can produce a system that is truly quiet and that does not impact performance.

Summary

There are many opportunities for PC motherboard manufacturers to differentiate their products based on the functionality that end users are looking for, and lowering acoustics is only one. As people spend more time with PCs, the ergonomic and environmental demands on these machines become more important. Today, there is an increasing awareness, among both PC users and manufacturers, that a low level of PC acoustic emissions must be regarded as an important factor in a comfortable working and home environment. Implementing continuous measurement and dynamic adaptation to temperature will make it possible to decrease the acoustic levels of PCs, to lower their power consumption, and to increase their reliability.

The ADM1027 *dB*COOL* controller, the ADP3163 multiphase VRM controller, and the AD1981A AC'97 SoundMAX* Codec are the result of a continuing commitment by Analog Devices to serve the PC market by providing products that monitor and manage power usage, process signals used in flat-panel displays and multimedia projectors, and enable PCs to provide CD-quality audio. ADI's SoundMAX audio software brings hyper-realistic audio functionality to PC motherboards and soundcards through the use of a variety of patented audio technologies. ADI also serves the high-end consumer market with integrated circuits used in products such as digital cameras and camcorders, DVD players, and surround-sound audio systems.

More Info

Analog Devices, Inc. is a leading manufacturer of precision, high-performance integrated circuits used in analog and digital signal processing applications. The company is headquartered in Norwood, Massachusetts, and has manufacturing facilities in Massachusetts, California, North Carolina, Ireland, the Philippines, Taiwan, and the United Kingdom. Analog Devices' stock is listed on the New York Stock Exchange, and the company is included in the S&P 500 Index. You can find out more by visiting their Web site.

References:

Bronzaft, A.L. & McCarthy, D.P. (1975). The effect of elevated train noise on reading ability. Environment and Behavior, 7, 517–528.
Bronzaft, A.L. (1981). The effect of a noise abatement program on reading ability. Journal of Environmental Psychology, 1, 215–222.
Bronzaft, A.L. (1996). Top of the Class. Norwood, NJ: Ablex.
Cohen, S., Evans, G.W., Krantz, & Stokols, D. (1980). Physiological, motivational, and cognitive effects of aircraft noise on children. American Psychologist, 35, 231–243.

Cohen, S., Glass, D.C. & Singer, J.D. (1973). Apartment noise, auditory discrimination and reading ability in children. Journal of Experimental and Social Psychology. 9, 407–422.
Collins, J. (February 3, 1997). "The Day Care Dilemma," Time.
Dallas, J.D. (1996). No more Jerichos! Hearing Rehabilitation Quarterly, 20, 9–11.
Deutsch, C.P. (1964). Auditory discrimination and learning: social factors. Merrill-Palmer Quarterly, 10, 277–296.
Evans, G.W. & Lepore, S.J. (1993). Nonauditory effects of noise on children: A critical review. Children's Environments, 10, 31–51.
Evans, G.W. & Lepore, S.J., (1993). Nonauditory effects of noise on children: A critical review. Children's Environments, 10(1), pp. 31–51.
Evans, G.W. & Maxwell, L., (1997). Chronic noise exposure and reading deficits: The mediating effects of language acquisition. Environment and Behavior, 29(5), pp. 638–656.
Evans, G.W., Hyge, S. & Bullinger, M. (1995). Chronic noise and psychological stress. Psychological Science, 6, 333–337.
Green, K.B., Pastenak, B.S. & Shore, R.E. (1982). Effects of aircraft noise on reading ability of school age children. Archives of Environmental Health, 37, 24–31.
Hambrick-Dixon, P.J. (1985). Effects of experimentally imposed noise on task performance of Black children attending day care center near elevated subway trains. Developmental Psychology, 22, 259–264.
Jones, F.N. & Tauscher, J. (1978). Residence under an airport landing pattern as a factor in teratism. Archives of Environmental Health, 33, 10–12.
Kryter, K.D. (1985). The effects of noise on man. Orlando, FL; Academic Press.
Kryter, K.D. (1985). The effects of noise on man. New York: Academic Press, Inc.
Stenzel, J. (1996). Flying off course: Environmental impacts of America's airports. New York: NY: Natural Resources Defense Council. Time (February 3, 1997). "How a child's brain develops: special report."
U.S. National Research Council (1982). Report of working group, 85, Prenatal effects of exposure to high-level noise. Washington, D.C.; National Academy Press.
Wachs, T. & Gruen, G. (1982). Early experience and human development. New York, NY, Plenum.

Networking and Communications

Tools for IOP321 with Intel® XScale™ Technology

Ed Whitty
Product Marketing Engineer
I/O and Bridges Division
Intel Corporation

Overview

Developers building storage applications based on the Intel® IOP321 processor have a number of advantages in designing top-of-the-line RAID on motherboard (ROMB), generic RAID adapters, host bus adapters, NAS (network attached storage), SAN (storage area network), and other storage applications products. These advantages include a PCI-X interface, single-chip integration of PCI, a combination of high-performance I/O application processing and extremely high data throughput, high-performance memory controllers, and integrated hardware-based Application Accelerator Units.

Developers basing new applications on the IOP321 also have a wide selection of tools and technologies to help them jump-start design, streamline development, and reduce time-to-market. For the first time, these tools and technologies were all demonstrated in a single location: the Intel Developer Forum (IDF) Conference Spring 2002.

In these demonstrations, developers learned about tools, technologies, and services that are optimized for the Intel® XScale™ microarchitecture and the 80321 product family. They met vendors offering technologies designed for use in concert with the Intel® IQ80321 reference platform; others provided their own choice of reference platforms, and still others offered porting services along with hardware and software products.

Essential to the Development of Embedded Storage Applications

Developers attending the demonstrations at IDF enjoyed a valuable first-hand look at some of the many products they depend on to get their IOP321-based products to market. This was the first time that third-party vendors offering products for storage-application developers gathered in one spot to showcase their tools.

For example, one vendor in this community, Wind River, has worked closely with Intel to offer a complete development solution for the IOP321: a real-time operating system, server I/O extensions, an optimized TCP/IP stack, JTAG hardware-assisted debugging, and a proprietary reference platform. Another vendor, Accelerated Technology Inc. (ATI), is offering a comprehensive development tools suite. This suite includes a robust integrated development environment (IDE) featuring a selection of tool chains from various vendors, among them ARM* and GNU, under a single interface. Known as the code|lab* Embedded Developers Suite, this product will soon ship with the Intel IQ80321 development platform.

All the tools and technologies shown by the vendors at IDF will play essential roles for storage-application developers. Most important, these tools play complementary roles with the IQ80321, and with one another, to provide developers of storage applications something they have not had before: a comprehensive IDE designed to help make the most of the high-performance/low-power-requirement advantages that the Intel® XScale™ microarchitecture and IOP321 are bringing to storage applications.

Tools for Diverse I/O Development Needs

Here are the vendors that were showcased at IDF, and some of the products they demonstrated for use in IOP321-based applications:

Accelerated Technology Inc. (ATI)

code|lab Embedded Developer Suite*, including design, development, and debug software; target-connection hardware and software; and evaluation boards or simulation software

Nucleus RTOS, including kernels; tools for prototyping, profiling, and Internet connectivity; Java* technologies; a file system; and a terminal application

ADI Engineering
80200EVB, a high-performance open reference design and evaluation platform for Intel XScale microarchitecture developers

Embedded Performance Inc. (EPI)
MAJIC Series*, multiprocessor advanced JTAG interface controllers (MAJIC) for a level of real-time debug visibility ordinarily associated with traditional in-circuit emulators and discrete microprocessors.

Macraigor Systems
*Raven**, a parallel port interface for JTAG or BDM debugging

Red Hat
*GNUPro**, embedded toolset and related services optimized for the Intel IQ80310 and IQ80321 Development Kits
*eCos**, an open-source real-time operating system for deeply embedded applications
*RedBoot**, a standardized embedded debug and bootstrap solution that provides firmware for running and debugging eCos, GNUPro applications, and embedded Linux* systems

TeamASA
NPWR, IOP321 and 82544 product board for NAS, RAID, and embedded Linux-based server solutions

Wasabi
NetBSD, a real-time operating system

Wind River Systems
*IxWorks**, a real-time operating system with server I/O extensions
Tornado for Intelligent Network Acceleration (TINA)*, a complete TCP/IP offload platform including VxWorks*, IxWorks*, and a proprietary reference platform
visionPROBE II, *visionICE II*, JTAG hardware-assisted debugging solutions for embedded systems, *IOP321 PCI I/O Card reference design*, a reference platform for the early prototyping of hardware and software

Summary

These demonstrations are further testimony to the growing support for the Intel XScale microarchitecture and the IOP321, and to the benefits of independent vendors working together toward a common goal: to promote and enable 80321 hardware and software design. It was also an opportunity for developers to get a hands-on demonstration of the tools and technologies that will help optimize development with the IOP321 processor and other Intel® I/O processors.

More Info

For more information on these vendors and others offering tools and technologies for developers of Intel IOP321-based storage solutions, visit the Intel® Intelligent Internet Storage Building Blocks Web site.

Author Bio

Ed Whitty is a product marketing engineer in the I/O and Bridges Division of Intel Corporation, where he is responsible for managing development tools supporting the Intel I/O processors and Intel XScale microarchitecture. He works with both Intel tool developers and developers from third-party software and hardware companies. Before joining Intel in 1997, Ed held positions with Sensormatic Electronics and served as an engineering officer in the United States Navy. Ed holds an M.B.A. from the University of Maryland at College Park and a B.S.M.E. from the University of Notre Dame.

Servers

UPnP Applications Enhance Mobile Functionality

Preston Hunt
Technical Marketing Manager, Network Architecture
Intel Labs
Intel Corporation

Ulhas Warrier
Engineering Manager, Network Architecture
Intel Labs
Intel Corporation

Overview

Just as the Internet has become an essential component of the mobile-computing experience, so too is another resource of data and information: the home network. With the advent of Universal Plug and Play* (UPnP*) technology and Internet gateway devices for the home, individuals can use smart phones, PDAs (personal digital assistants), and other mobile devices to tap into the home network through the Internet and thereby access their PC, audio/video equipment, appliances, and other networked devices securely and conveniently.

At the same time, the near market saturation of the current generation of mobile devices makes the addition of "secure home access" and related capabilities essential for manufacturers, opening new opportunities for software developers to build the capabilities into their mobile applications.

To support developers in this endeavor, this article provides a high-level overview of the following:
Secure home access and other capabilities that developers can incorporate into mobile-device applications with the help of UPnP technology
UPnP architecture
Tools and support organizations available to developers seeking to incorporate UPnP technology into their products.

The Value of Secure Home Access

The Internet is like a public library. It's an ideal place to store and access vast quantities of information with wide appeal—from weather reports to academic archives—but it's not the best repository for information with limited appeal, especially information of a private nature. Personal videos, photos, documents, and music or confidential information such as financial and tax data are more likely to be stored, and stored more securely, on a home PC. Being able to access such information over a cell phone or other mobile device is a valuable advantage for consumers.

For example, a business traveler is waiting for a flight at the airport when she spots a friend whom she hasn't seen in two years. Using her Pocket PC, she connects securely into her home computer through the UPnP IGD (Internet Gateway Device) that resides there and accesses selected photos for transfer to her Pocket PC so she can show them to her friend. Similarly, if this business traveler should find that her layover in a given city will be longer than scheduled, leaving unexpected time for a personal visit with a friend who lives in that city but who has an unlisted phone number, she can access the number stored on her home PC.

UPnP-enabled secure home access can also be used "in the other direction" for remote transmission of instructions to a variety of home-based appliances and other networked devices. For example, a hiker on a three-day mountain trip receives an alert on his Pocket PC that a front-porch security camera on his home network has detected motion. He securely accesses the network and downloads a live video snapshot from the camera to see a delivery person with a package the hiker has been expecting. The hiker then sends a message from his Pocket PC through the home-based IGD to a networked garage-door opener to enable temporary access for the delivery.

Related UPnP Capabilities Also Vital

Another application of UPnP technology is what's known as the personal Internet gateway. Like its IGD counterpart on the home network, a personal Internet gateway allows multiple devices in a personal area network to share a single Internet connection. These devices need not be Internet-capable but merely "network aware," which means they can access another device for the purpose of connecting to the Internet.

The advantage is that these devices can access the Internet without the costly software, hardware, and account services that a direct Internet connection requires. For example, just as she leaves the city, a vacationer realizes that she forgot to send an important e-mail. Instead of having to interrupt her vacation to return to the office, to seek out an Internet cafe, or to find a wireless LAN (802.11a/b) public-access point for access by her Pocket PC, the vacationer uses her cell phone as a personal Internet gateway to transmit the e-mail from her Pocket PC to the Internet.

Yet another UPnP application, proximity networking, provides connectivity among all devices within a certain context, such as a home, office, car, or a public space. For example, an individual sitting in the family room directs a UPnP-enabled Pocket PC to access music, photos, or videos stored on a PC in the home office for playback on an entertainment center in the family room. Similarly, a business traveler at the airport who is carrying a UPnP-enabled Pocket PC makes last-minute changes to a presentation he has stored there, remotely locates an available printing service in the airport, uploads the presentation, and has five copies printed for pick-up before boarding his flight.

Understanding UPnP Architecture

In practical terms, UPnP technology is a framework for defining the communication protocols between controllers, or control points, and devices. By freeing developers from having to design and debug systems to handle these activities, UPnP technology enables them to focus their efforts on their products and gain time-to-market advantages over competitors.

UPnP technology is broad in scope in that it supports not only home networks, but also proximity networks and networks in small businesses and commercial buildings. UPnP technology is powerful in that it enables data communication among any devices under the command of any control device on the network. UPnP technology is also architecturally neutral in that it works with any of the major operating systems (including Windows* and Linux*), programming languages (including C++* and Java*), processors (including IA32, Intel® XScale™, and several embedded controllers), and platforms.

Summary

The maturation of the mobile-phone market, the introduction of Internet gateway devices for the home, and the development of Universal Plug and Play (UPnP) technology make it imperative and easier than ever for software developers to enhance the connectivity and functionality of mobile applications. Secure home access, the personal Internet gateway, and proximity networking are a few of the powerful UPnP-based capabilities that can be incorporated into applications for mobile devices. Intel is strongly supporting developers pursuing this endeavor by providing a comprehensive software development kit (SDK), and participating in an active industry organization.

The open-source Intel® UPnP SDK is designed specifically to enable creators of home-networked devices and applications to add UPnP capabilities to their products. The kit provides a powerful application programming interface (API), source-code examples, and complete documentation for the implementation of UPnP-compliant control points and devices. Developers can download the SDK by joining the Intel® PCA Developer Network, an organization supporting the Intel® PCA open architectural framework, which separates development of applications and communication subsystems for faster development and deployment of wireless products and services.

Another organization, the UPnP Forum is devoted to the support of UPnP technology across the industry. Intel is a founding member of this consortium, which consists of more than 450 member companies—3Com, Alcatel, Cisco, Dell, Ericsson, HP, Maytag, Nokia, and Xerox, to name a few—from a variety of consumer industries.

More Info

Read the Universal Plug and Play white paper: "Using UPnP* technology to Extend the Reach of Handheld Devices." or for more information on UPnP architecture and related technologies, visit the following Web sites:
The Intel PCA Developer Network
Universal Plug and Play SDK
Wireless WAN

Author Bios

Preston Hunt is a technical marketing manager at Intel Labs, where he is focused on network architecture and trusted computing platforms for the home, including Universal Plug and Play, wireless networking, residential gateways, and end-to-end network security. Preston graduated summa cum laude from the Georgia Institute of Technology with a B.S.E.E. He currently has several patents pending on various home networking-related technologies and is an active participant in industry standards groups.

Ulhas Warrier is an architect and engineering manager in Intel Labs. He has been with Intel for over 10 years, and leads a group that develops technologies related to broadband gateway and home networking. Ulhas has been actively involved in the specification definition in the UPnP Forum Internet Gateway Working Committee since its inception. Ulhas has previously worked on IP telephony, video conferencing, and VPN technologies. He holds an M.S. in computer science.

—End of Intel Developer Update Magazine Issue 31—